

THE COMPLETE GUIDE TO: **MACHINE LEARNING** **IN THE ENTERPRISE**



Credit: iStock

Intelligent design



Machine learning is one of the premier buzzwords in not just the technology sector today, but across industry verticals. Before we dig into how machine learning algorithms are making businesses smarter than ever before though, let's attempt to tie down a working definition for the technology.

It is a subset of the umbrella technology discipline known as artificial intelligence (AI). The idea behind all of this technology is to grant computers a degree of independent thought and problem-solving skill, without the need for a developer or engineer to step in and tweak any code.

This is achieved by giving the machine large volumes of data that an algorithm can process and then learn from in order to make predictions and decisions for which it hasn't been specifically programmed.

This is slightly different to deep learning – another popular term in the world of AI – which is a type of machine learning inspired by the connections between neurons in the human brain. Here, researchers develop a man-made imitation of this biological connectivity known as artificial neural networks (known as neural nets).

When it comes to the enterprise, machine learning is an established technology, finally catching up after its early heyday, when IBM's Deep Blue mastered chess in 1997, firmly putting the technology on the map.

Most early use cases have been rooted in creating recommendation engines and personalizing e-commerce offerings, as pioneered and perfected by the likes of Amazon and, as you will see later on in this guide, travel


metasearch engine Expedia. However, using machine learning for fraud detection, cybersecurity and even trading on the stock market has also proved useful for large financial institutions.

Now the technology is starting to hit the mainstream, as cloud technology providers such as Amazon Web Services, Microsoft Azure and Google Cloud look to make it easier for businesses to adopt machine learning capabilities in the cloud, without having to find their own experts and build their own solutions from scratch.

We run through some of the best tools for businesses to look at when attempting to adopt machine learning, how vendors like Google are bringing the technology into the mainstream, and some examples of UK enterprises using machine learning to optimize existing business processes, or even create whole new lines of business off the back of the technology. [Scott Carey](#)

Contents

- 4** Best tools for developers and data scientists
- 9** Google's plans for AI and machine learning
- 13** Machine learning in cybersecurity
- 22** UK banks looking to use AI and machine learning
- 28** How Bloomberg hooks users to its terminals
- 33** How Expedia.com was built on machine learning



Credit: iStock

Best tools for developers and data scientists

Growing numbers of businesses are looking at how they can use machine learning in their operations

With businesses increasingly keen on incorporating artificial intelligence into their operations, machine learning – the ability for a system to learn from large data sets rather than following preset rules – offers a number of benefits. This might mean building predictive models for fraud prevention

in financial services, for example, or retailers making better recommendations to their customers.

Google, Microsoft, IBM and Amazon all offer machine learning APIs via their respective cloud platforms, making it easier for developers to build services by abstracting some of the complexity of their algorithms. There are also a growing number of open-source deep learning frameworks for data scientists.

Here are some of the top machine learning tools.

1. Azure Machine Learning Workbench

Microsoft announced a revamp of its Azure machine learning tools during its recent 'Ignite' conference. The firm revealed three major machine learning tools, one of which is the Azure Machine Learning Workbench, described as a cross-platform client for data and experiment management. The workbench will support modelling in Python, Scala and PySpark.

2. Azure Machine Learning Model Management

At its Ignite conference, Microsoft also unveiled the Azure Machine Learning Model Management tool. This aims to help developers 'manage and deploy machine learning workflows and models' while offering these modelling capabilities:

- Model versioning
- Model checking
- Deploying models to production
- Creating Docker containers with the models and testing them locally
- Automated model retraining
- Capturing model telemetry for actionable insights

3. Google APIs

Google has a host of machine learning tools on its Cloud Platform. These include the firm's popular Prediction API, which allows users to tap the search giant's algorithms to analyse data and predict future outcomes. The company has added further APIs to allow users to build their own machine learning-based services, including Speech, Translate and Vision.

In March 2017, Google launched a new machine learning API for automatically recognizing objects in videos and making them searchable. This API is called Cloud Video Intelligence and is to be used to help developers extract certain objects from videos automatically. In essence, the API allows developers to tag images in videos based on searchable terms for example, tree or house. Currently, developers can sign up for its beta version at tinyurl.com/nqyoze.

4. Amazon's Deep Scalable Sparse Tensor Network Engine (DSSTNE)

The open-source deep learning library, pronounced 'destiny', allows data scientists to train and deploy deep neural networks using GPUs. It can be seen as a response to Google's open sourcing of TensorFlow.

DSSTNE was built by the retail giant's engineers to power its recommendations engine, which makes product suggestions to the hundreds of millions of customers on its websites each day.

Amazon says: "We are releasing DSSTNE as open-source software so that the promise of deep learning can extend beyond speech and language understanding and object recognition to other areas such as search and recommendations. We hope that researchers around the

world can collaborate to improve it. But more importantly, we hope that it spurs innovation in many more areas.”

5. Amazon Web Services Machine Learning API

Amazon Web Services (AWS) launched its Amazon Machine Learning service in Europe August 2015, with the aim of making it easier for developers of all skill levels to access complex algorithms. The service was built on the technology used by its own internal data scientists. Amazon’s subsidiary says its machine learning service can generate billions of predictions a day, tapping into AWS data from services such as RedShift, S3 and its Relational Database Service.

6. Google TensorFlow

Google also open sourced its TensorFlow software library through an Apache licence, which powers many of its own services, including Google Photos, to Google Cloud Speech and is now being used by its DeepMind division for research. It can produce C++ or Python graphs that can be processed on CPUs or GPUs. These flow graphs depict the movement of data running through a system.

7. Microsoft Distributed Machine Learning Toolkit (DMLT)

Microsoft’s machine learning toolkit – which is available on GitHub – aims to ease crowded machine learning clusters, making it easier to run multiple (and differing) machine learning applications at the same time.

“Bigger models tend to generate better accuracies in various applications,” Microsoft says. “However, it


remains a challenge for common machine learning researchers and practitioners to learn big models.”

8. Microsoft Computational Network Toolkit (CNTK)

Also from Microsoft, the Computational Network Toolkit enables users to create neural networks depicted in directed graphs. While primarily made for speech recognition technology, since April 2015 it has become a more general machine learning toolkit supporting image, text and recurrent neural network (RNN) training.

9. IBM Watson Analytics

The Watson Analytics cloud service was unveiled in 2014 as part of IBM’s plans to turn Watson from a part-time game show contestant into a bona fide enterprise software proposition. It aims to help organizations that have little or no experience of predictive analytics put their business data to good use. IBM had already launched its Watson Developer Cloud – in 2013 – offering access to APIs via its Bluemix platform as a service cloud, allowing developers to create their own applications based on Watson’s smarts. [Christina Mercer](#)



Credit: Google

Google's plans for AI and machine learning

“Our goal is to create applications that can see, hear and understand,” says Jeff Dean, senior fellow at Google

Artificial intelligence is a clear priority for Google. In 2016, its AlphaGo program defeated world Go champion Lee Sedol, while its DeepMind arm is considered to be at the forefront of deep learning research. Machine learning also underpins many of its well-known products, from YouTube to Google Translate.

And now the firm is increasingly opening up its technology to other businesses as it attempts to grow its cloud operations.

At its Cloud Next conference Google unveiled the Cloud ML service, which helps machine learning engineers build “sophisticated, large” models based on its TensorFlow deep learning library, which was opened sourced in 2015.

The tech giant also added to its list of ‘pretrained’ machine learning platforms aimed at programmers without data science skills. Cloud Speech allows developers to convert audio to text by “applying powerful neural network models in an easy to use API”, the company says, and joins its Translate and Cloud Vision APIs.

How can businesses use machine learning?

Speaking to press at Google’s cloud user conference, Dean said that its machine learning tools will “allow other companies to build the same kind of understanding and insight from their data” that Google has achieved internally. This could be applied in retail, for instance.

“For example, if you have some data about customers and transactions and you want to predict something about it, that is a very general setting that can be used in all kinds of different enterprises,” Dean added. “So you could predict which customers are going to buy more than a hundred dollars of stuff and send them an offer.”

Online retail logistics firm Ocado Technology already uses machine learning algorithms in various ways across its business, such as powering its robotic systems. General manager James Donkin, says the company has been evaluating TensorFlow to help it build new algorithms, and would welcome the ability to move workloads to the cloud. “We really like the fact that we can run it locally, but if we want to use it on a larger

data volume it is easy to move it to the Google cloud to run the same TensorFlow code,” he reveals, adding that Google’s decision to open source the framework is “a big plus”. “We don’t like black box technologies that we can’t look inside.”

Democratizing machine learning access

Dean says that while the TensorFlow and Cloud ML tools are aimed at more sophisticated machine learning use cases, the cloud-based APIs cater to a wider audience of developers. One example is online hosting firm, Wix, which is using the Cloud Vision API. The technology is based on the image recognition that allows Google Photos to categorize photos of mountains or beaches, for example.

Dean explains: “Wix allows people to easily use websites and clients upload imagery. They want to be able to understand what kind of imagery it is, so that they can suggest the most appropriate content of that website.” He adds that there are also some that “can make pretty good use of the already trained APIs that we have released”.

Ocado’s Donkin says that the ease of use of the pretrained APIs means the firm could extend access to machine learning tools from its team of data scientists to developers. “One of our goals is to move beyond just data science using machine learning,” he told us. “So some of the new Google APIs – where normal developers can use machine learning – [could be] something we would do. We are looking at the structure we need to do that. So you have PhD specialist groups and you have software engineers, and how we can organizationally let that knowledge spread out will be.”

Future plans

Like many of Google's cloud services – such as BigQuery and Hadoop-based Dataproc – its machine learning services originate from technologies that it builds to run its massive operations. This includes Maps, Gmail and Android, as well as its robotics research.

“Machine learning is now used in lots and lots of Google products,” says Dean. “Some under the covers, so YouTube has lots of machine learning in it, but it is all in the recommendation engines.”

Google plans to continue to offer up its AI tools to its enterprise customers. “We think there is a lot of opportunity to introduce more and more of these pretrained APIs, and example models for particular kinds of settings. So recommendations might be a particular type of model,” Dean reveals.

That said, there are areas which Google is unlikely to make available to a wider audience, in particular the algorithms helping power its search business.

“It is generally a business decision about which ones make sense to release,” Dean explains. “We have been pretty open about the ones we have released. I would say things like search ranking we will probably not release but other than that most things [we would consider releasing]. We want people to use this to build their own cool products as well.” **Matthew Finnegan**



Credit: iStock

Machine learning in cybersecurity

How is machine learning being applied in cybersecurity and what does your organization need to know?

Recent breakthroughs in machine learning and artificial intelligence mean AI-enabled technologies are gaining traction. The billion-dollar cybersecurity industry is no exception, as vendors begin to scale and automate their processes intelligently – all while locked into the early stages of a security arms race with professional hackers.

A recent report from analyst firm ABI Research estimates that machine learning in cybersecurity will enormously bolster spending in big data, intelligence and analytics, reaching as much as \$96 billion (£71.9 billion) by 2021.

Vendors are likely to find buyers in large enterprises, and more than likely, across industries that are especially prone to attack: think government and defence, banking, and across the technology sector. At the moment, ABI's report says, User and Entity Behavioural Analytics – using machine learning for threat detection by analysing data at scale – is the driving force.

“Using static machine learning models to detect previously unknown malware is the only use case I'm aware of that offers clear evidence of effective results,” says cybersecurity analyst at 451 Research, Adrian Sanabria.

“Most machine learning use in the industry right now is experimentation and seeing what sticks,” he says. “The fact that machine learning has had some success in one area of infosec practically guarantees that this industry will attempt to use machine learning anywhere and everywhere it can be shoehorned in.”

But threat detection is not a trivial matter: in Cisco's recent annual cybersecurity report, it noted that the vast majority of companies are working to improve their threat detection capabilities.

There are plenty of public breaches where not only was the organization unaware of the intrusion until it was far too late, they had no idea about the true extent of the breach. A case in point is the devastating Yahoo hack – where eventually the company discovered all of its 3 billion email account details were compromised.

The author of the ABI paper, Dimitrios Pavlakis, says that to understand why machine learning is useful for detection, it's important to define the two primary distinctions for machine learning.

Supervised applications of machine learning tend to mean that you have clean and structured data – for example, anything that could be read in Excel – where you treat the model with what you know and what you then expect the software to do. In this case, you tell the algorithm what to do and where to look.

But unsupervised applications can examine unstructured data from multiple data sources. “With unsupervised models in machine learning you can then teach a model using neural networks and deep learning,” reveals Pavlakis. “You can teach a machine learning algorithm to detect the unknown. The algorithms are being trained, models are being stacked and trained all together.

“So you feed them data and tell them this could be normal, for example, and if something strange happens, the unseen threat can then be flagged.”

Some cybersecurity vendors in the machine learning space include Splunk, Gurucul, and Vectra, Trend Micro, Symantec, Invincea and CrowdStrike, and giant enterprises such as IBM are also doing work in the field.

Many eyes

“What machine learning offers is the ‘many-eyes’ option,” explains Gunter Ollmann, CSO at Vectra Networks. “You can use machines to observe the network continuously in real time, and correlate that across hundreds of millions, to trillions, of events on a daily basis.

“A traditional approach from a security practitioner perspective is to take logs, drop them into some central database, and then, offline, mine that data for events that we have a feeling might be there,” he says. “What machine learning offers is that all of the work can be done in real time, live in a network wire and without that human oversight.”

Andrew Gardner, senior director for machine learning at Symantec, explains that where machine learning will really help is in scale and automation. Think of the difference, he says, between two humans playing chess and two computers playing chess. And the computers can play each other at very high speeds.

“One thing that’s useful for is it allows us to do predictive testing,” he reveals. “We can, in a sandbox, use AI machine learning in the same way that an attacker might do, to predict and explore possible exploits on a scale that humans just can’t achieve.”

Today, machine learning is best suited for solving problems external to customers, says 451’s Adrian Sanabria. “Web reputation and malware prevention, for example, involve analysing publicly accessible data, whereas any use of machine learning internally by a customer may require extensive customization and training before it can be useful or effective,” he says.

That’s not to say internal applications will not be useful – but it is probably the sole domain of large enterprise-level organizations for now.

Chess

If threat detection is the most pertinent use of machine learning technologies today, what might this level of intelligence be able to achieve tomorrow? Developments

that are on the horizon suggest machine learning could soon help with other security applications.

“When repeated actions can be automated in a trusted way, we’ll see machine learning making incremental steps moving from detection, into remediation and then mitigation,” says Vectra’s Gunter Ollman. “A key part about this is that they are learning systems,” he reveals, adding that when Vectra’s systems are deployed in a live environment they never stop learning. They observe how the end user, the customer security expert behind that pane of glass is responding to things that are there.” Picture someone on the security team receiving a threat rating of 80 percent, and reacting by rating it much higher – the systems begin to learn the human processes of the organization.

Symantec’s Andrew Gardner adds that augmenting human talent will be a major part of the “natural evolution” of machine learning. “We have an email product and one of the key features people are interested in now is finding targeted spearphishing attacks,” he says.

“These mostly come through emails – it’s a tricky problem, machine learning wise, because a really well-targeted spearphishing attack is crafted programmatically at a level that’s hard to detect. It’s targeted to the right person at the right time, it probably doesn’t spam everybody and it may even be unique.

“We have analysts that can review suspicious emails, but without machine learning they are going through them one to ten at a time,” he continues. “We’ve built tools where you meld the human and the computer together, and the machine learning system learns to hunt for the person. It’s saying: ‘I can’t detect the threat, but

I can detect what the user looks for – I can cut through this in vast swathes of the time.’”

Gardner says without hyperbole that the developments we’ve started to see in machine learning are the first steps towards a self-evolution of technology – although he stresses that he doesn’t mean that “quite like a *Terminator* movie”.

“We have a system in our back end that teaches one of our antivirus detection engines how to improve itself,” Gardner says. “The engines are set up in an automated fashion, so it learns: ‘hey, you’ve made a mistake from this kind of data, you need to be focusing on this,’ etc.”

Rogue AI

The tendency for attackers and vendors to develop their capabilities in parallel complicates matters somewhat. If these incredible technologies can be used to defend, they can also be used to attack. It’s uncontroversial to say that vendors and attackers are locked into an arms race with no real end in sight.

According to ABI’s Dimitrios Pavlakis, the flock towards adopting machine learning is evidence of a reaction against the increasing sophistication and scale of attackers. “Machine learning in cybersecurity is, of course, a natural evolution of the technology,” he says. “But it is also a reactionary measure. Attackers by definition are always on the offensive, and they often have access to the same products that users and companies do – AV, security protocols, IP systems – they can study these inside and out.

“So attackers train algorithms to break into these products, they can buy an AV system, and they can test them – until their final product could be an algorithm

that can break the antivirus, and until the final product is ready to work on a larger percentage of machines and systems out there. Companies now make use of more sophisticated machine learning models as a result of higher threat factors.”

Splunk specializes in user behaviour analytics for threat detection. Matthias Maier, security evangelist for Splunk, predicts a more forbidding future for threats that make use of machine learning, especially as machine learning models become more accessible.

“What we expect will come up – though it’s not there yet – is that attackers might use machine learning to automate very targeted attacks,” he says. “Today they run a lot of complex attacks manually, and once they start using machine learning, they’ll be able to merge attacks like social engineering, research, phishing, delivery, credentials theft, and ransomware payouts. Once they start using machine learning to connect those different attack types they can automate manual processes.”

Vectra’s Gunter Ollman warns that professional attackers are studying machine learning very closely – and many of them are already data scientists.

“This is no different from 10 years ago when behavioural learning systems came out that the bad guys invested their own time, and they found ways to detect and bypass the sandboxing technologies,” he says. “I expect we’ll see that same level of thought and actions going into machine learning and artificial intelligence.”

Ollman adds attackers are already automating parts of their large-scale offensives. Worse still, there’s every possibility that malware equipped with AI could be set loose online, a rogue, intelligent design that

automatically and silently infiltrates systems for data. Traditionally, Ollman explains, the command and control server – used to remotely send malicious commands to botnets or other compromised systems – has been the “Achilles heel” for attacks.

“The area we’re most worried about is that many of today’s malware threat detection systems are focused on the command and control aspects of malware, VPNs, and other devices that have been installed inside corporate networks,” he says.

“The scary part is that as AI and machine learning advances, it is inevitable that these learning modes will be planted inside the malware. Once the malware has got inside the network, it will do away with the command and control necessity, and it will be automatically intelligent, and programmed to hunt out and seek data. The only time network traffic will be observed is when it’s completed the infiltration of that data.”

Snake oil

So it’s not all hype. Machine learning is concretely being used to protect enterprise-grade infrastructure today. But as with the development of any new technologies there is always a danger some vendors will latch on to the buzzword. What can organizations do to avoid buying into snake oil?

“I fundamentally believe that one of the best ways, and the best vehicles for understanding the scope and the capability of a technology from a vendor, is to look closely at their security research team and data science team,” says Vectra’s Gunter Ollmann. “First of all, if they don’t have either of those, then there is no way they can be doing machine learning or artificial intelligence.

“The second one is that effectively all the successful technologies that are out there in this space are products that are based on the capabilities of their research team – and the product is that research team wrapped into code, or into tin.

“So there are no security products out there that are more than what those research teams are capable of putting into that. Particularly when looking at the startup world, but even some of the larger vendors in the security space, a closer look at their security research teams and science teams, what their pedigree and what the sizes are, has proven to be a clear indicator of their capabilities, the product performance, and their ability to detect and mitigate a threat.”

It’s important, then, to have layers upon layer of protection, prevention and detection, says 451’s Adrian Sanabria – and not to be suckered in by slick marketing that paints machine learning as an all-curing panacea.

“We know from experience that attacks will simulate what infosec vendors are doing,” he argues. “I wouldn’t be surprised if they’ve already duplicated the industry’s machine learning work, and are working to determine ways to get around it, if they haven’t already.

“Machine learning models depend on a degree of likeness, so if attackers find a way to produce malware that looks significantly different from what models expect, machine learning-based detection methods could become ineffective overnight. This is one of the reasons it is important to have many layers of prevention, detection and hardening.” [Tamlin Magee](#)



Credit: iStock

UK banks looking to use AI and machine learning

UK banks are looking at the cutting-edge technology for everything from detecting fraud to building chatbots

The major UK banks are eyeing artificial intelligence (AI) technology to help them use the huge volumes of data they have on hand to improve compliance, increase customer engagement and improve operational efficiency.

Whether this truly benefits the customer more or the banks themselves is up for debate, and up to regulators

to try and police. Earlier this year, the father of the Internet Tim Berners-Lee warned against the possibility of AI systems becoming embedded into the financial world, and what that could mean for the fairness of the system. So just how are the major banks looking to use the cutting-edge AI and machine learning technologies? Here are just a few examples.

1. Anti-fraud

One of the core uses for machine learning in the banking world has been to combat fraud and improve compliance. The technology is ideally suited to the problem as machine learning algorithms can comb through huge transactional data sets to spot unusual behaviour.

“When you know about [fraud] now, something can be done about it,” Andrew McCall, chief engineer for big data at Lloyds Banking Group said earlier this year. “If you know about something that happened yesterday, it is not as effective as an anti-fraud mechanism.”

Douglas Flint, chairman of HSBC argued at the inaugural International FinTech Conference in April: “Using AI and machine learning to police the financial system is creating opportunities to do things better, to protect customers and ourselves.”

2. Algorithmic trading

Banks have used computer algorithms to trade stocks and shares since the 1970s, a practice that was partly responsible for the 1987 ‘Black Monday’ stock market crash. But as AI systems get better, the banks are naturally looking towards the technology to get an edge over the competition.

Michael Harte, head of group innovation at Barclays, said in April that the most obvious use for AI in banking is “in large algorithmic trading”. This means “using vast amounts of high velocity data to outsmart the competition and to provide better instruments and value to customers”.

It is worth noting that where Harte said customers, he presumably means Barclays’ investment clients, rather than everyday consumer banking customers. For more on electronic high frequency trading and how it helped a handful of banks ‘rig’ the US equity market Michael Lewis’ book *Flash Boys* is well worth a read.

3. Real time transaction analysis

Being able to track transactions in real-time has historically been an issue for major banks that have a huge amount of legacy IT infrastructure. However, getting data in place to be able to track transactions at low-latency would not only give the banks a better view of their customers, it would also give them the data set required to apply AI and deep learning to provide personalized, value added products to customers as it learns about spending habits over time.

The data science team at Lloyds has been working on this problem recently and has found a way to track transactions in near to real time. Andrew McCall, chief engineer for big data at Lloyds Banking Group said that getting to near real-time data processing within the bank “starts to open up lots of possibilities in terms of machine learning and how we can better serve customers and give them better insight into their own finances.”

4. Anti-money laundering

After receiving a \$1.9bn (£1.2bn) fine over money laundering in 2012, HSBC admitted that its controls were not fit for purpose. The bank said earlier this year that it is using Google Cloud machine learning capabilities for anti-money laundering.

Darryl West, CIO at HSBC, said the bank is using machine learning to run “analytics over this huge data set with great compute capability to identify patterns in the data to bring out what looks like nefarious activity within our customer base. The patterns that we identify are then escalated to the agencies and we work with them to track down the bad guys”. As startups such as the UK-based ComplyAdvantage try to show, AI is ripe for application to tracking money laundering as it is especially good at spotting odd behaviour within large data sets, such as banking transactions.

5. Personalized recommendations

Barclays head of group innovation Michael Harte envisions an AI system that can help the bank create and recommend better banking products to customers on a more personalized basis.

He said that the people within banks responsible for “inventing and maintaining products and services for customers” should be most excited by the technology. “So that what you bring customers is what they need. So tailor specific algorithms to the individual, instead of selling these generic products,” he added.

For example, there is already a growing number of start-ups that are looking to use machine learning algorithms to help customers find the best mortgage product, such as Trussle and Habito in the UK.

6. Credit applications

Banks can use machine learning algorithms to analyse an applicant for credit, be that an individual or a business, and make approvals according to a set of predefined parameters. These algorithms simply look at a customer's credit score, age and postcode to make a decision in seconds.

However, there are concerns that algorithms may not be as impartial as people think, and that getting the banks to explain the decision-making of their AI may not be as simple as it seems.

“It's not entirely clear how to properly equip a watchdog to do the job, simply because we are often talking about very complex systems that are unpredictable, change over time and are difficult to understand, even for the teams developing them,” Brent Mittelstadt of the Oxford Internet Institute told *The Guardian*.

7. Chatbots

Royal Bank of Scotland (RBS) has designed a customer service chatbot called Luvo, which is due to launch to customers in 2017 following a trial last year.

Luvo is a natural language processing AI bot, which will answer RBS, NatWest and Ulster bank customer's questions and perform simple banking tasks like money transfers. If Luvo is unable to find the answer it will pass a customer over to a member of staff. The bank said Luvo “talks to you through WhatsApp-type interaction” and what sets it apart from digital assistants like Siri and IKEA's Ask Anna is its ability to understand context and perform tasks.

8. Investment research

Outside of the UK, the Swiss bank UBS is developing an AI tool for investment research. The world's biggest wealth manager is building 'virtual agents' that can perform investment research to near-human levels.

Annika Schröder, AI lead at UBS Group Innovation, said that the bank is "trying to build virtual agents that can imitate the quality of an investment analyst. It can screen through market data, through SEC filings and can actually do a company valuation with all of the inputs that a human analyst would use and can produce text in a fairly decent quality and almost human mimicking language, so we are getting very close there."

UBS isn't ready to put this into live production yet, and Schröder wouldn't be pushed on a potential timeline. The bank is also investigating ways to use AI systems to eliminate bias from the behaviour of its portfolio managers; automate some KYC processes; customer-facing chatbots and automating the resolution of IT tickets. **Scott Carey**



Credit: iStock

How Bloomberg hooks users to its terminals

The head of data science at Bloomberg explains how the organization is increasingly turning to machine learning

Financial data specialist Bloomberg employs hundreds of data scientists to keep users hooked on its ubiquitous terminals – the keyboards and monitors that give financial staff access to reams of market information.

In the background, the Bloomberg Terminal crunches through 80,000 news wires, 4,000 FX feeds and 370

exchanges, driving around 60 billion data points a day to keep terminal users up to date on the global financial markets. The task for the data scientists at Bloomberg is to make this ocean of information discoverable and relevant for all of its 325,000 subscribers.

Gideon Mann is head of data science at Bloomberg. His role is to manage all of the data science that occurs across the large and varied organization.

“The bulk of data science that happens is building products,” he told us. “So my role ends up being mostly managing strategic, technical initiatives in basically three areas: natural language processing, search and machine learning that are embedded into products which serve the terminal.”

The Bloomberg Terminal

While many will think of Bloomberg as a media organization, subscriptions to the terminal and all of the associated data services – bundled together for around £20,000 a year – is used by thousands of bankers, traders, analysts and financial reporters, and is core to Bloomberg as a business. In terms of where the data science comes in, Bloomberg started by experimenting with machine learning for sentiment analysis nearly a decade ago. Mann admits that it took some time to get the organization to fully commit to machine learning – the computer science technique of teaching a machine to learn and adapt on the fly as it is fed large volumes of data – but the success of this project legitimized it for upper management.

“It took a number of years before the company realised that this particular competency takes a while,” he said. “Engineers can do it, but it is not simple. So then

the company started to commit to hiring and investing in quantitative programmers.” Now Bloomberg has between one and two hundred data science specialists within the organization, according to Mann.

Once the firm had proved the usefulness of the technique, and built the skills in-house, it started to apply the technique to the internal search on the terminal to improve data discovery through better ranking algorithms. A more recent project used computer vision to pick out data from tables embedded deep within financial reports and filings, a task that was once performed manually by programmers.

“What is much more successful is to use object recognition techniques on those tables,” Mann explained. “So it will recognize the boundary of the table, overlay with pick out columns and rows from that table into our database.” This means higher accuracy and speed.

Next, Mann wants to use techniques such as computer vision and natural language processing to improve the breadth of financial information available through the terminal. The aim is to allow users to increasingly make queries on the terminal using natural language instead of specialized commands.

“A lot of financial data is numbers, but a lot of the things that happen in the world that are pertinent to finance are expressed in language, either news stories we generate or aggregate, or press releases or documents the companies put out themselves, or even statement by officials,” Mann said. “All of that has a dramatic and fast change on the market. So the bulk of the data science and machine learning work we do is language processing, applying structure on top of it.”

Hiring

Mann believes that Bloomberg has got much better at hiring data scientists over the years, as it has grown to understand which people the organization needs – mainly computer science PhDs, if you were wondering. “I don’t want to come off as too braggadocious, but we have got better at [hiring]. We spend a lot of energy on it,” he explained. “We understand what we want and look for, and over the past year we have significantly increased the quality of applicants. They were always very good but the change has been we are able to hire people with a skill mix which is closer to what we need, cutting down on the training they need. I think we understand more of the people we need and the places we need to go and the universities they are coming out of.”

In essence, Bloomberg has steadily shifted away from statisticians and more towards quantitative programmers for any data science that occurs within its walls.

Mann is taking this strategy a step further. “We used to have the idea that each of these quantitative programmers had to be full stack,” he explained. “So take the data, clean it, structure it, support infrastructure, build a machine learning model, deploy it, babysit it and do fixes.” Now he wants smaller teams of specialists working on projects. For example, a data engineer, data scientist and production engineer working on a specific product within the terminal.

He recognizes that the industry is changing so fast that close ties to academia are integral to stay up to date with the latest technology trends.

Mann not only spends a lot of his time engaging with academia, either through publications, bringing in guest

speakers on a monthly basis or Bloomberg's own faculty grant programme, but the company itself promotes attending conferences for its technical staff.

"We send a tremendous amount of people to academic conferences, with the main aim being for them to learn and to be challenged by the experience of what is happening in academia," he said. For example, Bloomberg registered for 44 staff members to attend the Machine Learning Symposium in New York recently.

Open Source

In terms of tooling, Bloomberg has steadily shifted away from proprietary systems and vendors for data collection, processing and search to more open-source solutions such as Apache Spark and Apache Solr.

Mann admits that moving from vendors and proprietary software was something of a culture shift.

"When people talk about free software they say 'it is not free like beer, it is free like puppies' because it needs a lot of love and care," he said, adding that the people at Bloomberg eventually saw the benefit of contributing to open source and the sense of control it brings.

"Open source has really changed the way that we do business," Mann explained. "Traditionally, we built most of our stuff from scratch, for example generations of database technology, which created speed and reliability constraints. With big data processing, over the past five to ten years, the impact of Hadoop and now Spark has given us a whole new set of tools, and we are investing heavily in both of those. There was a time we were involved heavily with HBase but we are very aggressive with Spark right now. I don't know if we are an early adopter but we are certainly all in." **Scott Carey**



Credit: iStock

How Expedia.com was built on machine learning

Travel search giant Expedia has been building its core business on machine learning for the best part of a decade

Expedia has grown far beyond a search engine for flights – it’s now the parent company of a dozen travel brands, including Trivago and Hotels.com – but according to VP of global product David Fleischman, machine learning has always been at the heart of the company’s operations.

The business of delivering quality flight search results is tough, and Fleischman describes it as an “unbounded

computer science problem”. The reason for this is because flight itineraries and schedules are constantly changing, and Expedia’s proprietary ‘best fare search’ (BFS) has to ‘learn’ and adapt all the time.

The extent of the problem can be summed up by one statistic. The average Expedia.com flight search will take three seconds to deliver results. In those three seconds you will see, on average, 16,000 flight options, in order of convenience or price or time.

One weekend, the team at Expedia let BFS run for two full days on a single query: a round trip between Seattle and Atlanta in the United States. When they got back the algorithm had delivered “quadrillions of results”, says Fleischman.

This algorithm is always being tested and tweaked by the machine learning team at Expedia. The data scientists will test the algorithm against a whole range of bias, such as towards a business traveller or a family.

“We test multiple versions of the algorithm against one another and tweak the tuning on that,” says Fleischman. “We try these bias against one another and look at the result sets and ask if you get a better set.”

How do they define success, then? “We test and look at metrics to see if people bought more flights because we gave them a better result,” Fleischman explains.

Other applications of machine learning at Expedia

Expedia also uses its significant in-house machine learning resource – 700 data scientists and counting – to create algorithms for detecting fraud.

The next big project for Fleischman and his team is in natural language processing. In essence, this comes

back to the core aim of Expedia, to deliver the right results for a query. These queries could be a set of criteria like dates, destination and a price range, or a natural language query such as ‘I want to fly to Nice on Friday for the weekend’.

That’s the next challenge. “The main goal is to answer a traveller’s question and we use machine learning to solve that discovery problem,” Fleischman says.

In a blog post, he detailed the importance of natural language queries for mobile: “We’ve been experimenting with [natural language processing] for the past few years; ever since we realised that the standard travel search framework doesn’t work as well on mobile devices.”

Then there is personalization. Could a machine learning algorithm learn your preferences and travel habits over time and cut steps out of the buying process?

Fleischman says that the business is running experiments with personalization. “Personally,” he sighs, “it is something we need to look into but it is not necessarily something we know is the right thing as people just want to browse. So we think more about customization and detailed acknowledgement of where you are in your thinking.”

Put another way, where you are in the buying process is more important to Expedia than whether you like an aisle seat, at least at this point.

Nurturing machine learning talent

When it comes to hiring data scientists, Expedia puts a premium on problem solving skills.

“The trick with machine learning is the people not the code, the important bit is to understand the consumer problem,” Fleischman explains.

This chimes with the perspective of Nuno Castro, director of data science at Expedia, who recently told us: “People who can understand the organization from a commercial perspective, and who create relevant relationships in the organization will be more successful.”

Expedia works in agile teams, typically made up of a technologist, a product manager, a user experience designer and, according to Fleischman, machine learning experts are part of these core problem solving teams. Expedia doesn't dictate which tools they use or how they approach the problem.

And as Castro told us previously: “You and your team will typically be set a high level objective for which you need to determine the best cause of action. Unlike other areas, there is no one recipe for data science. Often you are not trying to find the right answer to a question, you're trying to find the right questions to ask in the first place.” **Scott Carey**



©2017 International Data Group.